Parallel mining of time–faded heavy hitters

Massimo Cafaro[a,*], Marco Pulimeno[a], Italo Epicoco[a]

[a] *University of Salento, Lecce, Italy*

arXiv:1701.03004v1 [cs.DS] 11 Jan 2017

Abstract

We present PFDCMSS, a novel message–passing based parallel algorithm for mining time–faded heavy hitters. The algorithm is a parallel version of the recently published FDCMSS sequential algorithm. We formally prove its correctness by showing that the underlying data structure, a sketch augmented with a Space Saving stream summary holding exactly two counters, is mergeable. Whilst mergeability of traditional sketches derives immediately from theory, we show that merging our augmented sketch is non trivial. Nonetheless, the resulting parallel algorithm is fast and simple to implement. To the best of our knowledge, PFDCMSS is the first parallel algorithm solving the problem of mining time–faded heavy hitters on message–passing parallel architectures. Extensive experimental results confirm that PFDCMSS retains the extreme accuracy and error bound provided by FDCMSS whilst providing excellent parallel scalability.

Keywords: message–passing, heavy hitters, time fading model, sketches.

1. Introduction

In this paper we deal with the problem of mining in parallel time–faded heavy hitters (also called frequent items), and we present PFDCMSS, a novel message–passing based parallel algorithm which is a parallel version of the recently published FDCMSS sequential algorithm [4].

Mining of heavy hitters in a data stream has been thoroughly studied, and the problem is regarded as one of the most important in the streaming algorithms literature. Depending on the particular application, the problem is reported in the literature as *hot list analysis* [19], market basket analysis [2] and *iceberg query* [17], [1].

Even though there are many possible applications, we recall here some of the most important contexts to which the problem has been successfully applied: network traffic analysis [14], [16], [27], analysis of web logs [7], Computational and theoretical Linguistics [18].

All of the algorithms for detecting heavy hitters can be classified as being either *counter* or *sketch* based, the difference being that counter–based algorithms rely on a set of counters which are used to keep track of stream items, whilst sketch–based algorithms monitor the data stream by using a sketch data structure, often a bi-dimensional array data structure containing a counter in each cell. Stream items are mapped by hash functions to corresponding cells in the sketch. The former algorithms (counter–based) are deterministic, whilst the latter (sketch–based) are probabilistic.

Regarding counter–based algorithms, the first sequential algorithm has been designed by Misra and Gries [26]. Their algorithm was rediscovered, independently, about twenty years later by Demaine et al. [14] (this algorithm is known in the literature as the *Frequent* algorithm) and Karp et al. [22]. Among the developed counters–based algorithms we recall here *Sticky Sampling* and *Lossy Counting* [24], and *Space Saving* [25]. Sketch–based solutions include *CountSketch* [7], *Group Test* [12], *Count-Min* [11] and *hCount* [21].

Relevant parallel algorithms include [6], [3] and [5] which are message-passing based parallel versions of the Frequent and Space Saving algorithms. Shared-memory algorithms have been designed as well, including a parallel version of Frequent [31], a parallel version of Lossy Counting [30], and parallel versions of Space Saving [28] [13].

*Corresponding author

Email addresses: massimo.cafaro@unisalento.it (Massimo Cafaro), marco.pulimeno@unisalento.it (Marco Pulimeno), italo.epicoco@unisalento.it (Italo Epicoco)

Recent shared-memory parallel algorithms for heavy hitters were recently proposed in [29]. Finally, accelerator based algorithms exploiting a GPU (Graphics Processing Unit) include [20] and [15]. Regarding related work, i.e., parallel algorithms specifically designed to solve the problem of mining time–faded heavy hitters, we are not aware of any other algorithm: to the best of our knowledge, ours is the first parallel algorithm solving the problem on message–passing parallel architectures.

In this paper, we are concerned with the problem of detecting in parallel heavy hitters in a stream with the additional constraint that recent items must be weighted more than former items. The underlying assumption is that, in some applications, recent data is certainly more useful and valuable than older, stale data. Therefore, each item in the stream has an associated timestamp that will be used to determine its weight. In practice, instead of estimating items' frequencies, we are required to estimate items' *decayed frequencies*.

This paper is organized as follows. We recall in Section 2 preliminary definitions and concepts that will be used in the rest of the manuscript. We present in Section 3 our PFDCMSS algorithm and formally prove in Section 4 its correctness. Next, we provide extensive experimental results in Section 5, showing that PFDCMSS retains the extreme accuracy and error bound provided by the sequential FDCMSS whilst providing excellent parallel scalability. Finally, we draw our conclusions in Section 6.

2. Preliminary definitions

In this Section we introduce preliminary definitions and the notation used throughout the paper. We deal with an input data stream σ consisting of a sequence of n items drawn from a universe $\mathcal{U}$; without loss of generality, let m be the number of distinct items in σ i.e., let $\mathcal{U} = \{1, 2, \ldots, m\}$, which we will also denote as $[m]$. Let f_i be the frequency of the item $i \in \mathcal{U}$ (i.e., its number of occurrences in σ), and denote the frequency vector by $\mathbf{f} = (f_1, \ldots, f_m)$. Moreover, let $0 < \phi < 1$ be a support threshold, $0 < \epsilon < 1$ a tolerance such that $\epsilon < \phi$ and denote the 1-norm of $\mathbf{f}$ (which represents the total number of occurrences of all of the stream items) by $\|\mathbf{f}\|_1$.

In this paper, we are concerned with the problem of detecting in parallel frequent items in a stream with the additional constraint that recent items must be weighted more than former items. The underlying assumption is that, in some applications, recent data is certainly more useful and valuable than older, stale data. Therefore, each item in the stream has an associated timestamp that will be used to determine its weight. In practice, instead of estimating frequencies, we are required to estimate *decayed frequencies*. Two different models have been proposed in the literature: the *sliding window* and the *time fading* model. PFDCMSS works in the latter model. Furthermore, even though the basic ideas underlining the algorithm are also appropriate for an online distributed setting, here we are assuming that the entire dataset is available for offline processing.

The time fading model [23] [9] [8] does not use a window sliding over time; freshness of more recent items is instead emphasized by *fading* the frequency count of older items. This is achieved by computing the item's *decayed frequency* through the use of a decay function that assign greater weight to more recent occurrences of an item than to older ones: the older an occurrences is, the lower its decayed weight.

Definition 1. Let $w(t_i, t)$ be a decayed function which computes the decayed weight at time t for the occurrence of item i arrived at time t_i. A decayed function must satisfy the following properties:

1. $w(t_i, t) = 1$ when $t_i = t$ and $0 \leq w(t_i, t) \leq 1$ for all $t > t_i$;
2. w is a monotone non-increasing function as time t increases, i.e., $t' \geq t \implies w(t_i, t') \leq w(t_i, t)$.

Related work has mostly exploited *backward decay* functions, in which the weight of an item is a function of its age, a, where the age at time $t > t_i$ is simply $a = t - t_i$. In this case, $w(t_i, t)$ is given by $w(t_i, t) = \frac{h(t-t_i)}{h(t-t)} = \frac{h(t-t_i)}{h(0)}$, where h is a positive monotone non-increasing function.
The term backward decay stems from the aim of measuring from the current time back to the item's timestamp. Prior algorithms and applications have been using backward exponential decay functions such as $h(a) = e^{-\lambda a}$, with $\lambda > 0$ as decaying factor.

In our algorithm, we use instead a forward decay function, defined as follows (see [10] for a detailed description of the forward decay approach). Under forward decay, the weight of an item is computed on the amount of time between the arrival of an item and a fixed point L, called the *landmark* time, which, by convention, is some time earlier than

the timestamps of all of the items. The idea is to look forward in time from the landmark to see an item, instead of looking backward from the current time.

Definition 2. Given a positive monotone non-decreasing function g, and a landmark time L, the forward decayed weight of an item i with arrival time $t_i > L$ measured at time $t \geq t_i$ is given by $w(t_i, t) = \frac{g(t_i - L)}{g(t - L)}$.

The denominator is used to normalize the decayed weight so that $w(t_i, t)$ is always less than or equal to 1 as requested by Definition 1.

Definition 3. The *decayed frequency* of an item v in the input stream σ, computed at time t, is given by the sum of the decayed weights of all the occurrences of v in σ: $f_v(t) = \sum_{v_i = v} w(t_i, t)$.

Definition 4. The *decayed count* at time t, $C(t)$, of a stream σ of n items is the sum of the decayed weights of all the items occurring in the stream: $C(t) = \sum_{i=1}^{n} w(t_i, t)$.

The Approximate Time–Faded Heavy Hitters (ATFHH) problem is formally stated as follows.

Problem 1. *Approximate Time–Faded Heavy Hitters. Given a stream σ of items with an associated timestamp, a threshold $0 < \phi < 1$ and a tolerance $0 < \epsilon < 1$ such that $\epsilon < \phi$, and letting g be a decaying function used to determine the decayed frequencies and t be the query time, return the set of items F, so that:*

- *F contains all of the items v with decayed frequency at time t $f_v(t) > \phi C(t)$ (decayed frequent items);*
- *F does not contain any item v such that $f_v(t) \leq (\phi - \epsilon)C(t)$.*

In the following, when clear from the context, the query time shall be considered an implicit parameter, so we write f_v and C instead of $f_v(t)$ and $C(t)$. The algorithm presented makes use of a Count–Min sketch data structure augmented by a Space Saving summary associated to each sketch cell. In the following, we recall the main properties of the Count–Min and the Space Saving algorithms in the case of non decaying frequencies, but the same properties also hold in a time-fading context.

Count–Min is based on a sketch whose dimensions are derived by the input parameters ϵ, the error, and δ, the probability of failure. In particular, for Count–Min $d = \lceil \ln 1/\delta \rceil$ is the number of rows in the sketch and $w = \lceil e/\epsilon \rceil$ is the number of columns. Every cell in the sketch is a counter, which is updated by hash functions. By using this data structure, the algorithm solves with probability greater than or equal to 1 - δ the *frequency estimation* problem for arbitrary items. The algorithm may also be extended to solve the *approximate frequent items* problem as well, by using an additional heap data structure which is updated each time a cell is updated. Since in Count-Min the frequencies stored in the cells overestimate the true frequencies, a point query for an arbitrary item simply inspects all of the d cells in which the item is mapped to by the corresponding hash functions and returns the minimum of those d counters.

Space Saving is a counter-based algorithms solving the heavy hitters problem. It makes use of a stream summary data structure composed by a given number of counters $k \ll n$, n being the length of the stream. Each counter monitors an item in the stream and tracks its frequency. A substitution strategy is used when the algorithm processes an item not already monitored and all of the counters are occupied.

Let σ be the input stream and denote by $\mathcal{S}$ the summary data structure of k counters used by the Space Saving algorithm. Moreover, denote by $|\mathcal{S}|$ the sum of the counters in $\mathcal{S}$, by f_v the exact frequency of an item v and by $\hat{f}_v$ its estimated frequency, let $\hat{f}^{min}$ be the minimum frequency in $\mathcal{S}$. If there exist at least one counter not monitoring any item, $\hat{f}^{min}$ is zero.

Finally, denote by $\mathbf{f} = (f_1, \ldots, f_m)$ the frequency vector. The following relations hold (as shown in [25]):

$$|\mathcal{S}| = \|\mathbf{f}\|_1, \tag{1}$$

$$\hat{f}_v - \hat{f}^{min} \leq f_v \leq \hat{f}_v, \qquad v \in \mathcal{S}, \tag{2}$$

$$f_v \leq \hat{f}^{min}, \qquad v \notin \mathcal{S}, \tag{3}$$

$$\hat{f}^{min} \leq \left\lfloor \frac{\|\mathbf{f}\|_1}{k} \right\rfloor. \tag{4}$$

Therefore, it holds that

$$\hat{f}_v - f_v \leq \hat{f}^{min} \leq \left\lfloor \frac{\|\mathbf{f}\|_1}{k} \right\rfloor, \quad v \in \mathcal{U}. \tag{5}$$

3. The algorithm

In this section, we start by recalling our sequential algorithm FDCMSS [4]. The key data structure is an augmented Count–Min sketch D, whose dimensions d (rows) and w (columns) are derived by the input parameters ϵ, the error, and δ, the probability of failure. Whilst every cell in an ordinary CM sketch contains a counter used for frequency estimation, in our case a cell holds a Space Saving stream summary with exactly two counters. The idea behind the augmented sketch is to monitor the time–faded items that the sketch hash functions map to the corresponding cells by an instance of Space Saving with two counters, so that for a given cell we are able to determine a *majority item candidate* with regard to the sub-stream of items falling in that cell.

Indeed, by using a data structure $\mathcal{S}$ with two counters in each cell, and letting $C_{i,j}$ denote the total decayed count of the items falling in the cell $D[i][j]$, the majority item is, if it exists, the item whose decayed frequency is greater than $\frac{C_{i,j}}{2}$. The corresponding majority item candidate in the cell is the item monitored by the Space Saving counter whose estimated decayed frequency is maximum. We have proved that, with high probability, if a time-faded item is frequent, then, in at least one of the sketch cells where it is mapped, it is a majority item with regard to the sub-stream of items falling in the same cell. Therefore, our algorithm will detect it.

Theorem 1. *If an item i is frequent, then it appears as a majority item candidate in at least one of the d cells in which it falls, with probability greater than or equal to* $1 - (\frac{1}{2\phi w})^d$.

Regarding the error bound of our algorithm, let f_i be the exact decayed frequency of item i in the stream σ and $\hat{f}_i$ be the estimated decayed frequency of item i returned by FDCMSS. Let C be the total decayed count of all of the items in the stream. We have proved the following error bound.

Theorem 2. *$\forall u \in [m]$, $\hat{f}_u$ estimates the exact decayed count f_u of u at query time with error less than ϵC and probability greater than $1 - \delta$.*

The proofs of aforementioned theorems can be found in [4].

The algorithm's initialization requires as input parameters ϵ, the error; δ, the probability of failure; and ϕ, the support threshold. The initialization returns a sketch D. The procedure starts deriving $d = \lceil \ln 1/\delta \rceil$, the number of rows in the sketch and $w = \lceil \frac{e}{2\epsilon} \rceil$, the number of columns in the sketch. Then, for each of the $d * w$ cells available in the sketch D we allocate a data structure $\mathcal{S}$ with two Space Saving counters c_1 and c_2. Given a counter c_j, $j = 1, 2$, we denote by $c_j.i$ and $c_j.f$ respectively the counter's item and its estimated decayed frequency. Finally, we set the support threshold to ϕ, select d pairwise independent hash functions $h_1, \ldots, h_d : [m] \to [w]$, mapping m distinct items into w cells, and initialize the *count* variable, representing the total decayed count of all of the items in the stream, to zero.

Updating the sketch upon arrival of a stream item i with timestamp t_i, shown in pseudo-code as Algorithm 1, requires computing x, which is the non normalized forward decayed weight of the item, and incrementing *count* by x. Then, we update the d cells in which the item is mapped to by the corresponding hash functions $h_j(x)$, $j = 1, \ldots, d$ by using the Space Saving item update procedure.

Let $\mathcal{S}$ denote the Space Saving stream summary data structure with two counters corresponding to the cell to be updated. Updating $\mathcal{S}$ upon arrival of an item works as follows. When processing an item which is already monitored by a counter, its estimated frequency is incremented by the non normalized weight x. When processing an item which is not already monitored by one of the available counters, there are two possibilities. If a counter is available, it will be in charge of monitoring the item, and its estimated frequency is set to the non normalized weight x. Otherwise, if all of the counters are already occupied (their frequencies are different from zero), the counter storing the item with minimum frequency is incremented by the non normalized weight x. Then, the monitored item is evicted from the

Algorithm 1 Process

Require: i, an item; t_i, timestamp of item i; D, sketch data structure
Ensure: update of sketch related to item i; update the local total decayed count.

```
1: procedure PROCESS(i, t_i, D)
2:     x ← g(t − t_i)                          ▷ compute the non normalized decayed weight of item i
3:     lCount ← lCount + x                     ▷ update local total decayed count
4:     for j = 1 to d do
5:         S ← D[j][h_j(i)]
6:         SpaceSavingUpdate(S, i, x)          ▷ update the sketch
7:     end for
8: end procedure
```

counter and replaced by the new item. This happens since an item which is not monitored can not have a frequency greater than the minimal frequency.

PFDCMSS, the parallel version of our sequential algorithm, works as follows. We assume the offline setting in which the stream items have been stored as a static dataset along with the corresponding timestamps. It is worth noting here immediately that our algorithm works in the streaming (online) setting as well. Indeed, in the former case (offline setting) we partition the input dataset and timestamps using a simple 1D block-based domain decomposition among the available p processes and then process in parallel the sub-streams assigned to the processes using Algorithm 1. In the latter case (online setting), we have instead p distributed sites, each handling a different stream $\sigma_i, i = 1, \ldots, p$ processed again using Algorithm 1.

In the parallel version, once the sub-streams have been processed, one of the processes is in charge of determining the time–faded heavy hitters. In order to do so, all of the processes engage in a parallel reduction in which their sketches are merged into a global sketch which preserves all of the information stored in the local sketches. This sketch is then queried and the time–faded heavy hitters are returned.

In the distributed setting, one of the sites may act as a centralized coordinator or there can be another different site taking this responsibility. The coordinator broadcasts, when required, a "query" message to the p sites, which then temporarily stop processing their sub-streams, and engage in the sketch merge procedure. We can imagine the distributed sites as being multi-threaded processes, in which one thread executes Algorithm 1, temporarily stops when a query message is received from the coordinator, creates a copy of its local sketch and then resume stream processing whilst another thread engages in the distributed sketch merging procedure using the sketch copy.

In order to retrieve the time–faded heavy hitters, a query can be posed when needed. The query, shown in pseudo-code as Algorithm 2, starts by determining the global decayed count for the whole stream σ. This requires a parallel reduction in which the local decayed counts are summed. It is worth noting here that the global decayed count is still non normalized; the normalization occurs dividing by $g(t - L)$, where t is the query time and L denotes the landmark time. Then, we build, through a user's defined parallel reduction, a global sketch G which is obtained by merging the local sketches. To do so, each process invokes a parallel reduction by using the MergeSketch operator shown in pseudo-code as Algorithm 3.

The sketches are reduced as follows: for every corresponding cell in two sketches to be merged, the hosted Space Saving summaries are merged following the steps described in [5], i.e., building a temporary summary $\mathcal{S}_C$ consisting of all of the items monitored by both $\mathcal{S}_1$ and $\mathcal{S}_2$. To each item in $\mathcal{S}_C$ is assigned a decayed frequency computed as follows: if an item is present in both $\mathcal{S}_1$ and $\mathcal{S}_2$, its frequency is the sum of the its corresponding frequencies in each summary; if the item is present only in one of either $\mathcal{S}_1$ or $\mathcal{S}_2$, its frequency is incremented by the minimum frequency of the other summary. At last, in order to derive the merged summary, we take only the 2 items in $\mathcal{S}_C$ with the greatest frequencies and discard the others.

It is worth noting here that the sum of the counters in the stream summary data structure $\mathcal{S}$ related to a given cell $D[i][j]$ is equal to the value that the Count–Min sketch–based algorithm would store in the counter variable corresponding to that cell, i.e., the 1-norm of the frequency vector corresponding to the sub–stream falling in the cell through the pairwise independent hash functions. Thus, an augmented sketch is equivalent, from this perspective, to a Count–Min sketch and this property is preserved by the merge procedure. From now on we will call this property

Algorithm 2 Query

Require: t, query time; D, process' sketch
Ensure: set of frequent items

```
1:  procedure QUERY(t)
2:      gCount ← ParallelReduction(lCount, Sum)
3:      gCount ← gCount / g(t−L)
4:      G ← ParallelReduction(D, MergeSketch)
5:      F = ∅
6:      for each S_ij ∈ G do
7:          let c_1 and c_2 be the counters in S_ij
8:          c_m ← argmax(c_1, c_2)                ▷ c_m the counter with maximum decayed count
9:          if c_m.f / g(t−L) > φ * gCount then
10:             p ← PointEstimate(c_m.i, t)
11:             if p > φ * gCount then
12:                 F ← F ∪ {(c_m.i, p)}
13:             end if
14:         end if
15:     end for
16:     return F
17: end procedure
```

1-*norm equivalence*.

However, merging Count–Min sketches simply requires adding the corresponding cells' counters. Indeed, via linearity, the sum of sketches is equal to the sketch of the sums. Instead, in our case, we need an ad hoc procedure in order to correctly merge the two Space Saving stream summaries hosted by the corresponding cells so that 1-norm equivalence property is preserved. Nonetheless, the augmented sketch which results from our parallel merge reduction is 1-norm equivalent to the Count–Min sketch obtained by summing the Count–Min sketches corresponding to our augmented sketches which are the input of the parallel merge reduction.

Once the global sketch is obtained, the query procedure initializes F, an empty set, and then it inspects each of the $d * w$ cells in the sketch D. For a given cell, we determine c_m, the counter in the data structure $\mathcal{S}$ with maximum decayed count. We normalize the decayed count stored in c_m dividing by $g(t-L)$, and then compare this quantity with the threshold given by $\phi * gCount$ ($gCount$ being the normalized global decayed count). If the normalized decayed frequency is greater, we pose a point query for the item $c_m.i$, shown in pseudo-code as Algorithm 4. If p, the returned value, is greater than the threshold $\phi * gCount$, then we insert in F the pair $(c_m.i, p)$.

The point query for an item j returns its estimated decayed frequency. After initializing the *answer* variable to infinity, we inspect each of the d cells in which the item is mapped to by the corresponding hash functions, to determine the minimum decayed frequency of the item. In each cell, if the item is stored by one of the Space Saving counters, we set *answer* to the minimum between *answer* and the corresponding counter's decayed frequency. Otherwise (none of the two counters monitors the item j), we set *answer* to the minimum between *answer* and the minimum decayed frequency stored in the counters. Since the frequencies stored in all of the counters of the sketch are not normalized, we return the normalized frequency *answer* dividing by $g(t-L)$.

At the end of the query procedure the set F is returned.

4. Correctness

Here, we prove that our algorithm correctly merges two FDCMSS sketches. The merge procedure preserves all of the properties of the sketch, including the fact that, considering the sum of the Space Saving counters in each sketch cell, an FDCMSS sketch is 1-norm equivalent to the classical Count–Min sketch.

It is worth noting here that we would obtain a correct result by using the merge procedure presented in [5] to combine the Space Saving summaries stored in the corresponding sketch cells, but we also want to impose 1-norm

Algorithm 3 MergeSketch

Require: D_1, D_2: sketches to be merged.
Ensure: G, the merged sketch

procedure MERGESKETCH(D_1, D_2)
  for each $\mathcal{S}^1_{ij} \in D_1, \mathcal{S}^2_{ij} \in D_2$ **do**
    $m_1 \leftarrow \text{MIN}(c^1_1.f, c^1_2.f)$ ▷ m_1, the minimum of counters' frequency in $\mathcal{S}^1_{ij}$
    $m_2 \leftarrow \text{MIN}(c^2_1.f, c^2_2.f)$ ▷ m_2, the minimum of counters' frequency in $\mathcal{S}^2_{ij}$
    for each $c_{s_1} \in \mathcal{S}_1$ **do**
      $c_{s_2} \leftarrow \text{FIND}(\mathcal{S}_2, c_{s_1}.i)$
      if c_{s_2} **then**
        $c_{s_c}.f \leftarrow c_{s_1}.f + c_{s_2}.f$
        DELETE($\mathcal{S}_2, c_{s_2}$)
      else
        $c_{s_c}.f \leftarrow c_{s_1}.f + m_2$
      end if
      $c_{s_c}.i \leftarrow c_{s_1}.i$
      INSERT($\mathcal{S}_C, c_{s_c}$)
    end for
    for each $c_{s_2} \in \mathcal{S}_2$ **do**
      $c_{s_c}.i \leftarrow c_{s_2}.i$
      $c_{s_c}.f \leftarrow c_{s_2}.f + m_1$
      INSERT($\mathcal{S}_C, c_{s_c}$)
    end for
    PURGE($\mathcal{S}_C$) ▷ $\mathcal{S}_C$ now contains 2 counters with the greatest frequencies
    $G[i][j] \leftarrow \mathcal{S}_C$
  end for
  return G
end procedure

Algorithm 4 PointEstimate

Require: j, an item; t, query time
Ensure: estimation of item j decayed count;

procedure POINTESTIMATE(j, t)
  $answer \leftarrow \infty$
  for $i = 1$ to d **do**
    $\mathcal{S} \leftarrow G[i][h_i(j)]$ ▷ let c_1 and c_2 be the counters in $\mathcal{S}$
    if $j == c_1.i$ **then**
      $answer \leftarrow \text{MIN}(answer, c_1.f)$
    else if $j == c_2.i$ **then**
      $answer \leftarrow \text{MIN}(answer, c_2.f)$
    else
      $m \leftarrow \text{MIN}(c_1.f, c_2.f)$
      $answer \leftarrow \text{MIN}(answer, m)$
    end if
  end for
  return $\frac{answer}{g(t-L)}$
end procedure

equivalence, i.e., the additional condition that the sum of counters' values in each merged cell always reflects the total decayed count of the items which fell in the corresponding cells.

Indeed, in [5] we showed how to merge Space Saving stream summaries in parallel. However, we have proved that our merge procedure satisfies the Space Saving properties described by eq. 2-5, and the following relaxed version of eq. 1:

$$|\mathcal{S}| \leq \|\mathbf{f}\|_1, \tag{6}$$

As shown in Theorem 3, which is the main result of this section, it turns out that $k = 2$ counters (i.e., majority item mining) is a special case: when the Space Saving summaries to be merged hold two counters, than the property in eq. 1 holds for the merged summary in its original form, that is $|\mathcal{S}| = \|\mathbf{f}\|_1$, without modifying the merge procedure designed in [5].

Theorem 3. *The parallel merge algorithm provides an augmented sketch that preserves all of the properties of a FDCMSS sketch.*

Proof. The correctness of the parallel FDCMSS sketch merge algorithm derives from the correctness of the Space Saving merge procedure, already shown in [5]. It remains to show that, when looking to the sum of the Space Saving counters associated to each cell, the merged augmented sketch is still 1-norm equivalent to a Count–Min sketch, that is, the sum of the counters values is equal to the decayed count of all the items fallen in that cell.

Let us recall the merge algorithm for Space Saving summaries introduced in [5]. We will use the multiset notation, thus let us rewrite the properties of a Space Saving summary stated in equations 1-4, this time with reference to multisets. Indeed, we model the input stream as a multiset (also called a *bag*), which essentially is a set where the duplication of elements is allowed. We shall use a calligraphic capital letter to denote a multiset, and the corresponding capital Greek letter to denote its *underlying* set. In particular, we extend the traditional notion of multiset as follows. Instead of considering an indicator function which returns the multiplicity of an item, we use a function providing the decayed frequency of that item. Therefore, summing over all of the items we obtain the total decayed count in place of the cardinality of the multiset.

Definition 5. A decayed multiset $\mathcal{N} = (N, f_{\mathcal{N}})$ is a pair where N is some set, called the underlying set of elements, and $f_{\mathcal{N}} : N \to \mathbb{R}$ is a function which provides the *decayed frequency* for each $x \in N$ according to Definition 3.

The decayed count of $\mathcal{N}$ is expressed by

$$|\mathcal{N}| := \sum_{x \in N} f_{\mathcal{N}}(x), \tag{7}$$

whilst the cardinality of the underlying set N is

$$|N| := Card(N) = \sum_{x \in N} 1. \tag{8}$$

From now on, when referring to either the exact or estimated frequency of an item, we shall mean the item's exact or estimated decayed frequency. Recall that our Space Saving stream summary data structure uses exactly $k = 2$ counters, and let $\mathcal{N} = (N, f_{\mathcal{N}})$ be the input decayed multiset, $\mathcal{S} = (\Sigma, \hat{f}_{\mathcal{S}})$ the decayed multiset of all of the monitored items and their respective counters at the end of the sequential Space Saving algorithm's execution, i.e., the algorithm's summary data structure. Let $|\mathcal{S}|$ be the sum of the frequencies stored in the counters, $f_{\mathcal{N}}(e)$ the exact frequency of an item e, $\hat{f}_{\mathcal{S}}(e)$ its estimated frequency and $\hat{f}_{\mathcal{S}}^{\,min}$ the minimum frequency in $\mathcal{S}$, where $\hat{f}_{\mathcal{S}}^{\,min} = 0$ when $|\Sigma| < 2$. Indeed, even though a summary data structure has exactly 2 counters, it may monitor less than 2 items, since an item is actually monitored if and only if its counter's frequency is different from zero. The following relations hold, for each item $e \in N$:

$$|\mathcal{S}| = |\mathcal{N}|, \tag{9}$$

$$\hat{f}_{\mathcal{S}}(e) - \hat{f}_{\mathcal{S}}^{\,min} \leq f_{\mathcal{N}}(e) \leq \hat{f}_{\mathcal{S}}(e), \qquad e \in \Sigma, \tag{10}$$

$$f_{\mathcal{N}}(e) \leq \hat{f}_{\mathcal{S}}^{\,min}, \qquad e \notin \Sigma, \tag{11}$$

$$\hat{f}_{\mathcal{S}}^{\,min} \leq \left\lfloor \frac{|\mathcal{N}|}{2} \right\rfloor. \tag{12}$$

Now, let $\mathcal{S}_1 = (\Sigma_1, \hat{f}_{\mathcal{S}_1})$ and $\mathcal{S}_2 = (\Sigma_2, \hat{f}_{\mathcal{S}_2})$ be two summaries related respectively to the input sub-arrays $\mathcal{N}_1 = (N_1, f_{\mathcal{N}_1})$ and $\mathcal{N}_2 = (N_2, f_{\mathcal{N}_2})$, with $\mathcal{N} = \mathcal{N}_1 \uplus \mathcal{N}_2 = (N, f_{\mathcal{N}})$. Let $\mathcal{S}_M = (\Sigma_M, \hat{f}_{\mathcal{S}_M})$ be the final merged summary.

Theorem 3 in [5] states that if eqs. (10) - (12) hold for $\mathcal{S}_1$ and $\mathcal{S}_2$ and, if it is verified a relaxed version of eq. (9), i.e., it holds that

$$|\mathcal{S}_i| \leq |\mathcal{N}_i|, \qquad i = 1, 2 \tag{13}$$

then these properties continue to be true also for $\mathcal{S}_M$ (it is worth noting here that eq. (13) also holds for summaries produced by the sequential Space Saving algorithm). The authors show that this is enough to guarantee the correctness of the merge operation, but, in general $|\mathcal{S}_M| \leq |\mathcal{N}|$.

In order to obtain $\mathcal{S}_M$, we start combining $\mathcal{S}_1$ and $\mathcal{S}_2$ to obtain $\mathcal{S}_C$, and then, if $|\Sigma_C| > 2$, we take the two counters with the greatest frequency values in $\mathcal{S}_C$ in order to build $\mathcal{S}_M$, otherwise we return $\mathcal{S}_M = \mathcal{S}_C$.

We can express the combine operation as shown by the following equation:

$$\hat{f}_{\mathcal{S}_C}(e) = \begin{cases} \hat{f}_{\mathcal{S}_1}(e) + \hat{f}_{\mathcal{S}_2}(e), & e \in \Sigma_1 \cap \Sigma_2, \\ \hat{f}_{\mathcal{S}_1}(e) + \hat{f}_{\mathcal{S}_2}^{min}, & e \in \Sigma_1 \setminus \Sigma_2, \\ \hat{f}_{\mathcal{S}_2}(e) + \hat{f}_{\mathcal{S}_1}^{min}, & e \in \Sigma_2 \setminus \Sigma_1, \end{cases} \tag{14}$$

In the special case of stream summaries holding exactly $k = 2$ counters, it holds that for $i = 1, 2$, $|\Sigma_i| \leq 2$, and $|\Sigma_C| \leq 4$. Now, suppose that $|\mathcal{S}_i| = |\mathcal{N}_i|$, (this is true when $\mathcal{S}_i$ is produced by the sequential Space Saving) and let $\delta = \hat{f}_{\mathcal{S}_1}^{min} + \hat{f}_{\mathcal{S}_2}^{min}$ and $x = |\Sigma_C| - 2$. Furthermore, suppose that the entries in $\mathcal{S}_C$ are sorted in ascending order with regard to the counters frequencies.

As proved in [5], it holds that:

$$|\mathcal{S}_C| = |\mathcal{S}_1| + |\mathcal{S}_2| + x\delta = |\mathcal{N}_1| + |\mathcal{N}_2| + x\delta \tag{15}$$

$$|\mathcal{S}_M| = |\mathcal{N}_1| + |\mathcal{N}_2| + x\delta - \sum_{i=1}^{x} \hat{f}_{\mathcal{S}_C}(e_i), \tag{16}$$

where the sum is extended over the first x entries.

We have to show that the difference $x\delta - \sum_{i=1}^{x} \hat{f}_{\mathcal{S}_C}(e_i)$ is always equal to zero when $k = 2$, so that $|\mathcal{S}_M| = |\mathcal{N}|$.

When $x \leq 0$, $x\delta = 0$. In that case, $\mathcal{S}_M = \mathcal{S}_C$ and $|\mathcal{S}_C| = |\mathcal{N}|$.

When $x > 0$, the first x counters of $\mathcal{S}_C$ have values equal to δ. To see this, consider the two cases $x = 1$ and $x = 2$.

When $x = 2$, that is, the two summaries to be merged contain different items and $|\Sigma_C| = 4$, this is easily seen by simple computations: in fact, δ is the minimum value a counter in $\mathcal{S}_C$ can assume, and there are at least two counters with this value in $\mathcal{S}_C$, obtained combining the two counters with minimum value in $\mathcal{S}_1$ and $\mathcal{S}_2$. As a consequence these counters are the first two, and $x\delta - \sum_{i=1}^{x} \hat{f}_{\mathcal{S}_C}(e_i) = 0$.

When $x = 1$, one of the following cases arises:

1. one of the summaries (without loss of generality, let us suppose it is $\mathcal{S}_1$) contains two counters, the other summary ($\mathcal{S}_2$) contains only one counter and no item is in common between the summaries. In this case, δ is equal to the minimum counter in $\mathcal{S}_1$ since $\hat{f}_{\mathcal{S}_2}^{min} = 0$, but it is also the minimum counter in $\mathcal{S}_C$, hence it holds that $\delta - \hat{f}_{\mathcal{S}_C}(e_1) = 0$

2. both summaries contain two counters and they have exactly an item in common. In this case we further have to distinguish three cases:
   (a) the item in common has the minimum frequency in both the summaries. The combined frequency of this item will be equal to δ which is the sum of the minimum frequencies of two summaries. Its combined frequency is also the minimum in $\mathcal{S}_C$, hence it holds that $\delta - \hat{f}_{\mathcal{S}_C}(e_1) = 0$;
   (b) the item in common has the maximum frequency in both the summaries. Its combined frequency is also the maximum value in $\mathcal{S}_C$, and $\mathcal{S}_C$ contains two distinct items with combined frequency equal to δ which is also the minimum in $\mathcal{S}_C$, hence $\delta - \hat{f}_{\mathcal{S}_C}(e_1) = 0$;
   (c) the item in common appears with minimum frequency in one summary (without loss of generality, let us suppose in $\mathcal{S}_1$) and with maximum frequency in the other summary ($\mathcal{S}_2$). The combined frequency of the item which appears with minimum frequency in $\mathcal{S}_2$ is equal to δ which again is the minimum frequency of the counters in $\mathcal{S}_C$, hence $\delta - \hat{f}_{\mathcal{S}_C}(e_1) = 0$.

Taking into account that in all of the cases when $x = 1$ the $\mathcal{S}_C$ contains at least one item whose combined frequency is equal to δ, it holds that $x\delta - \sum_{i=1}^{x} \hat{f}_{\mathcal{S}_C}(e_i) = 0$.

□

We have shown that all of the properties of a Space Saving summary of two counters are preserved by the merge procedure introduced in [5]. It suffices to guarantee that all of the properties stated for an FDCMSS sketch continue to hold after the parallel merge procedure depicted in the algorithm presented. In particular, it holds the property 1, which guarantees that a merged FDCMSS sketch continues to be 1-norm equivalent to a Count–Min sketch.

5. Experimental results

In this section, we report experimental results on synthetic datasets. Here, we thoroughly test our algorithm using an exponential decay function. All of the experiments have been carried out on the Galileo cluster machine kindly provided by CINECA in Italy. This machine is a linux CentOS 7.0 NeXtScale cluster with 516 compute nodes; each node is equipped with 2 2.40 GHz octa-core Intel Xeon CPUs E5-2630 v3, 128 GB RAM and 2 16 GB Intel Xeon Phi 7120P accelerators (available on 384 nodes only). High-Performance networking among the nodes is provided by Intel QDR (40Gb/s) Infiniband. All of the codes were compiled using the Intel C++ compiler v17.0.0.

Let f be the true frequency of an item and $\hat{f}$ the corresponding frequency reported by an algorithm, then the Relative Error is defined as $\Delta f = \frac{|f-\hat{f}|}{f}$, and the Average Relative Error is derived by averaging the Relative Errors over all of the measured frequencies.

Precision, a metric defined as the total number of true heavy hitters reported over the total number of candidate items, quantifies the number of false positives reported by an algorithm in the output stream summary. Recall is the total number of true heavy hitters reported over the number of true heavy hitters given by an exact algorithm. In all of the results we obtained 100% recall, even on a tiny sketch of size 4 x 800 (recall may be less than 100%, but this happens only when the sketch size is really minimal). For this reason, to avoid wasting space, we do not show here recall plots. Rather, we present Precision, Absolute Error, Average Relative Error (ARE), Updates/ms and runtime/performance plots since we are interested in understanding the error behavior and the algorithm's scalability when we use an increasing number of cores of execution. Table 1 reports the experiments carried out. For each different metric under examination, we varied n, the stream size in billions of items, ρ, the skew of the zipfian distribution, ϕ, the threshold and w, the number of sketch columns. All of the other parameters are fixed when varying one of the previous ones, and we show, on top of each plot, the fixed parameters' values.

Finally, we also present, for the metrics of interest, the results obtained by fixing the stream size and varying the number of cores utilized from 1 to 512. We conclude this section with a comparison between strong and weak scalability.

With the experiment 1 we aim at measuring the algorithm accuracy, the experiment 2 aims at measuring how the parallelization affects the algorithm's accuracy, finally experiment 3 is meant to measure the computational performance of the parallel algorithm measuring both strong and weak scalability.

Table 1: Design of experiments. The input stream size n is expressed in billions, ρ denotes the skewness of the input data distribution, ϕ is the support threshold, w represents the number of columns in the sketch data structure, p is the number of cores and gs denotes the grain size, i.e. the number of elements (expressed in billions) in the input stream for each core. The *nf* value stands for *not fixed*

Exp.	Aim	Varying	n	ρ	ϕ	w	p	gs
1	Algorithm accuracy	$n = \{1, 2, 4, 8\}$	nf	1.1	0.01	1340	16	0.5
		$\rho = \{1, 1.4, 1.8, 2.2\}$	8	nf	0.01	1340	16	0.5
		$\phi = \{0.001, 0.004, 0.008, 0.016\}$	8	1.1	nf	1340	16	0.5
		$w = \{800, 1600, 3200, 6400\}$	8	1.1	0.01	nf	16	0.5
2	Parallel alg. accuracy	$p = \{1, 16, 128, 256, 512\}$	8	1.1	0.01	1340	nf	nf
3	Computational performance	$p = \{1, 16, 128, 256, 512\}$	8	1.1	0.01	1340	nf	nf
		$p = \{1, 16, 128, 256, 512\}$	nf	1.1	0.01	1340	nf	0.5

5.1. Algorithm accuracy

As shown by Figure 1, our parallel algorithm provides 100% Precision in all of the experiments carried out. Both the Absolute and the Average Relative Error, depicted respectively in Figure 2 and 3, have extremely low values, in particular with regard to their mean values. We observe that the Absolute Error is only slightly affected by the stream size n, and it's not affected at all by the threshold ϕ. The behaviour observed when varying ρ and w is expected. We observe a decrement of the Absolute Error in both cases, since, when the skew is higher, the number of frequent items in the corresponding zipfian distribution is lower and, when w is higher, increasing the sketch size provides better accuracy and, correspondingly, less error. Regarding the Average Relative Error, we observe the same qualitative behaviors in the experiments carried out.

Finally, the updates done per millisecond, shown in Figure 4, appear to be stable around 100,000 when varying the stream size n and the threshold ϕ. There is a visible increment (from 100,000 to 120,000) when varying ρ and a decrement (from 100,000 to 90,000) when varying w. These behaviors are expected for the same reasons we gave when analyzing the error. Indeed, processing a stream in which the number of frequent items is lower is usually faster. On the other hand, increasing the sketch size provides better accuracy, but more time is required to update the sketch.

5.2. Impact of the parallelization on the accuray

We now discuss the experimental scalability shown by our parallel algorithm. Figure 5 provides the results for the metrics under examination when testing strong scalability. That is, we fix the problem size (i.e., the stream size n) and increase the number of cores on which the algorithm is executed. As shown, Precision, Absolute and Average Relative Error are not affected at all by a strong scaling of the application, with Precision always equal to 100% and extremely low error values. Finally, the observed increment of the Updates/ms when varying the number of cores utilized is expected, due to the frequency updates made in parallel. Ideally, the throughput of the algorithm, measured as updates/ms, should increase with the same rate of increase of the number of MPI processes.

5.3. Computational perfomance

Figure 6 is related to the comparison we did to test weak scalability, which refers to the scalability of a parallel application obtained when the problem size is increased along with the number of cores, so that we measure how the running time changes with regard to the number of cores for a fixed problem size per core (whilst, for strong scalability, we measure how the running time changes with regard to the number of cores for a fixed total problem size).

As shown, the plot for strong scaling is a log-log plot of the running time versus the number of cores. The dashed straight line with slope -1 indicates ideal scalability, whereas any upward curvature away from that line indicates limited scalability. The plot reports a good strong scalability even on 512 cores; this is due to the high number of items to be processed in the input stream which makes the computational time higher than the parallel overhead.

Regarding weak scalability, the corresponding plot provides an indication of loss of performance when scaling from 1 to 16 cores while we have a very good scalability from 16 cores up to 512. This can be explained considering

the parallel architecture used for testing. One computing node is made of two octa-core Xeon processors, so that the cores share the main memory banks and the third level cache memory. In the weak scalability experiment the problem size increases linearly with the number of processes, hence also the total memory increases linearly; since from 1 to 16 cores we use only one computing node the memory contention between parallel processes increases. When varying from 16 to 512 cores we use several different computing nodes ranging from 1 to 32; each node runs 16 processes which compete for memory accesses as already discussed, hence the further slight loss of performance is due to the communication overhead.

6. Conclusions

We have presented PFDCMSS, a novel message–passing based parallel algorithm for mining time–faded heavy hitters, which, to the best of our knowledge, is the first parallel algorithm solving the problem on message–passing parallel architectures. We have formally proved its correctness by showing that the underlying data structure, is non trivially mergeable. However, the parallel algorithm is fast and simple to implement, and we have shown, through extensive experimental results, that PFDCMSS retains the extreme accuracy and error bound provided by FDCMSS whilst providing very good parallel scalability.

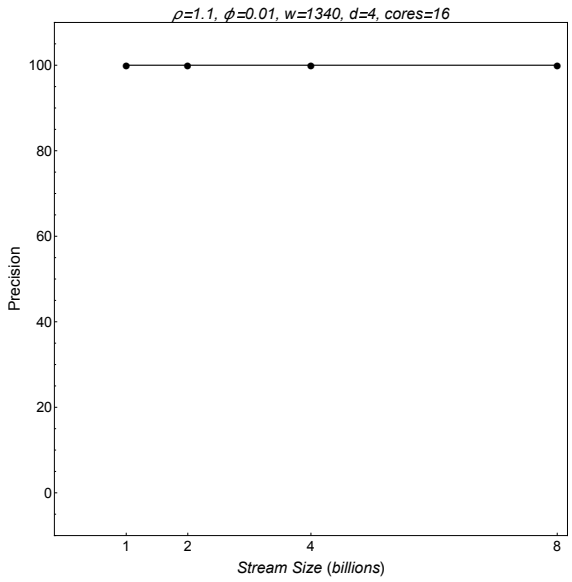

(a) varying n

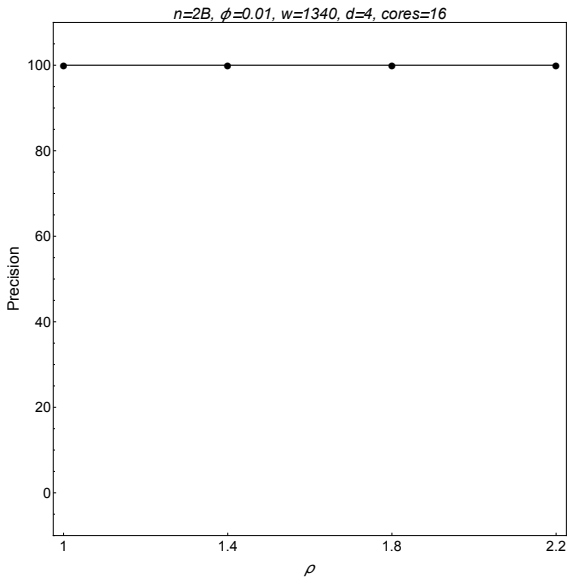

(b) varying ρ

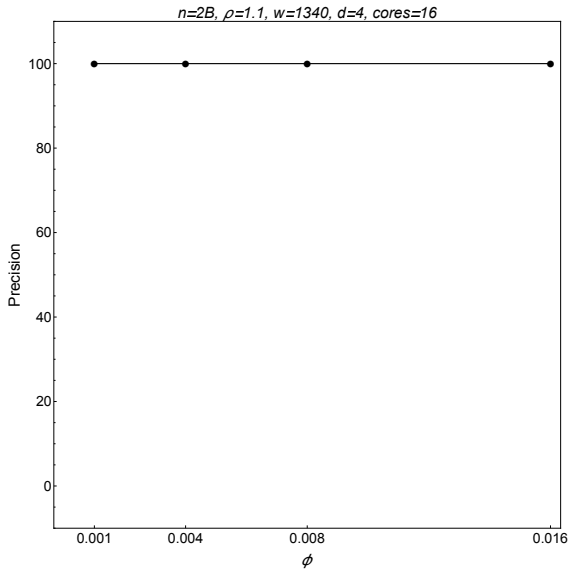

(c) varying ϕ

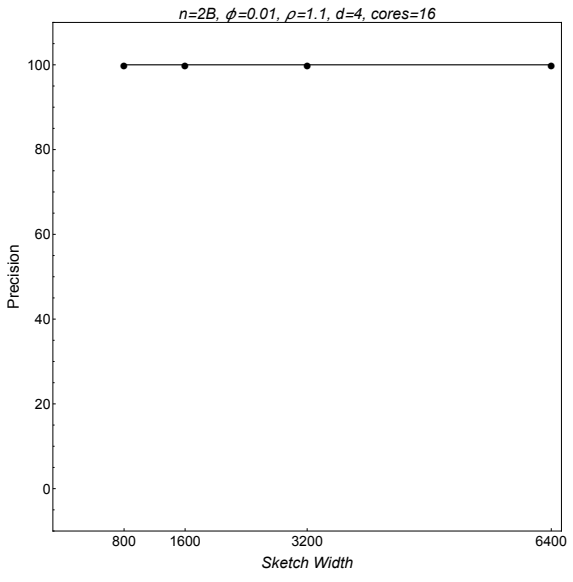

(d) varying w

Figure 1: Precision

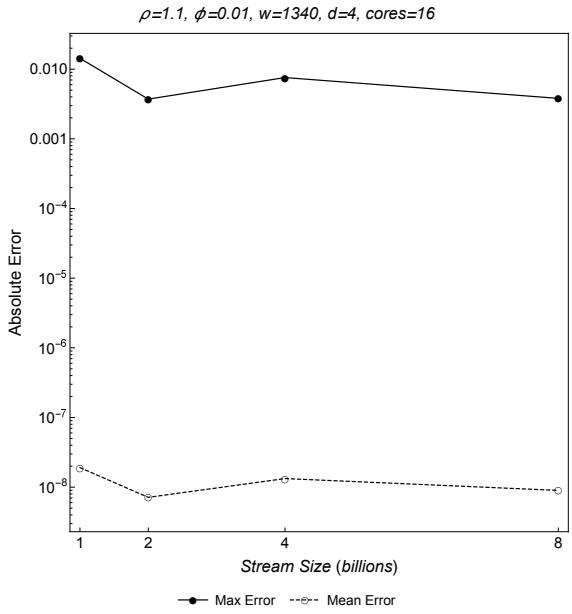

(a) varying n

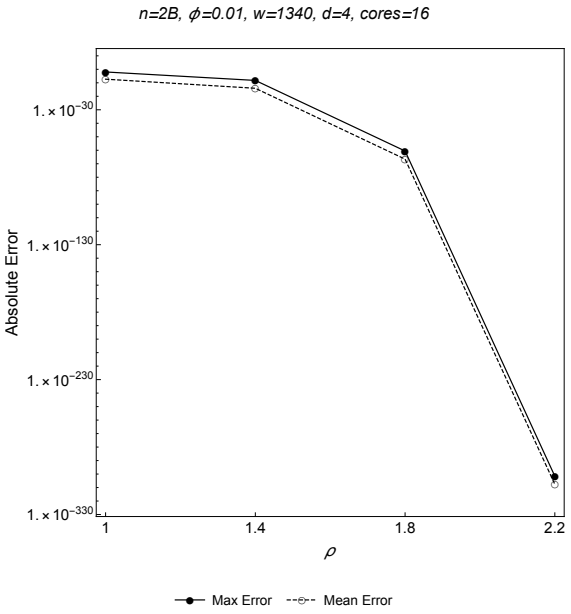

(b) varying ρ

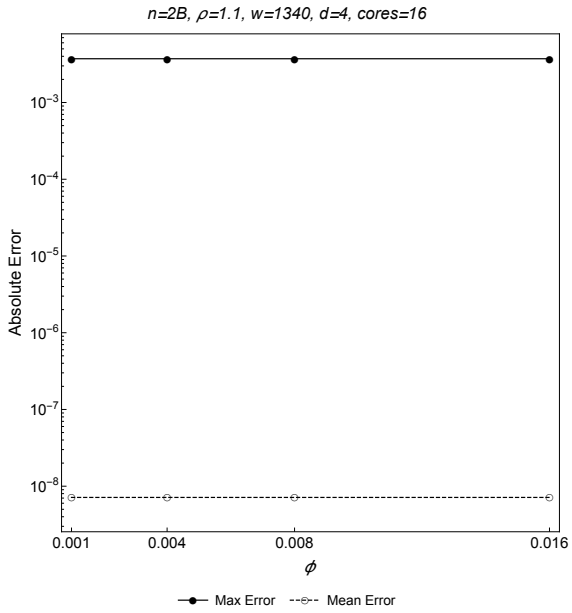

(c) varying ϕ

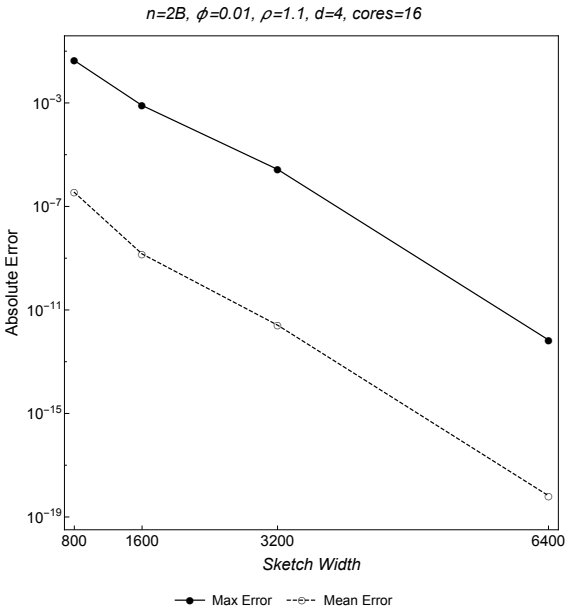

(d) varying w

Figure 2: Absolute Error

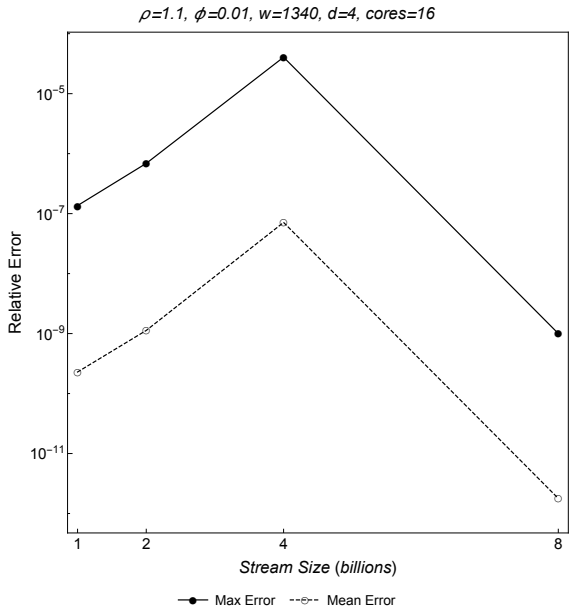

(a) varying n

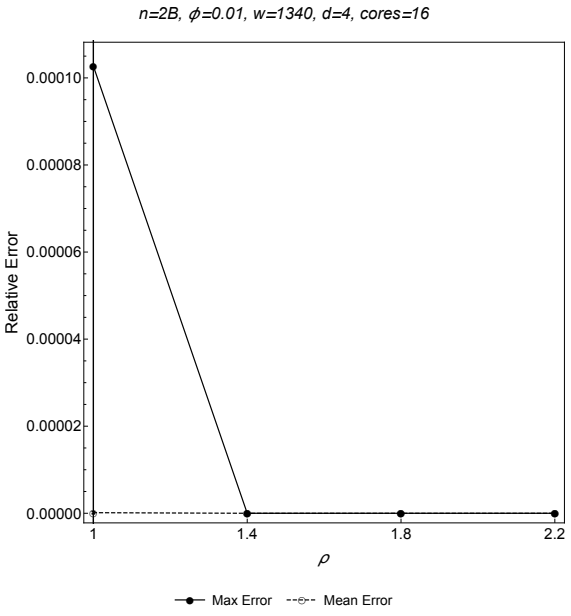

(b) varying ρ

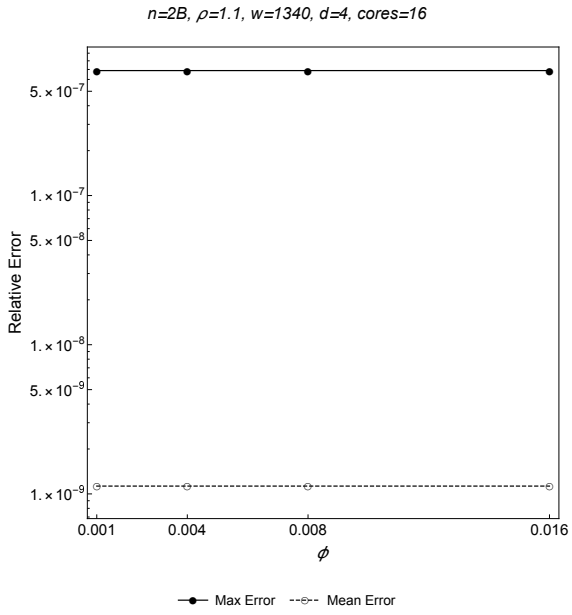

(c) varying ϕ

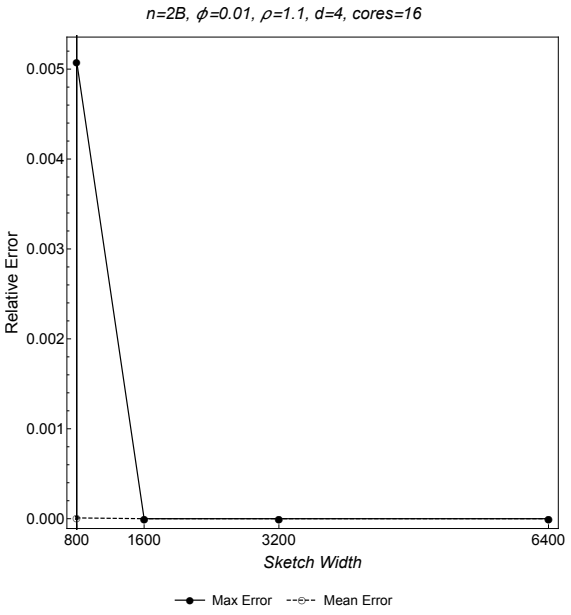

(d) varying w

Figure 3: Relative Error

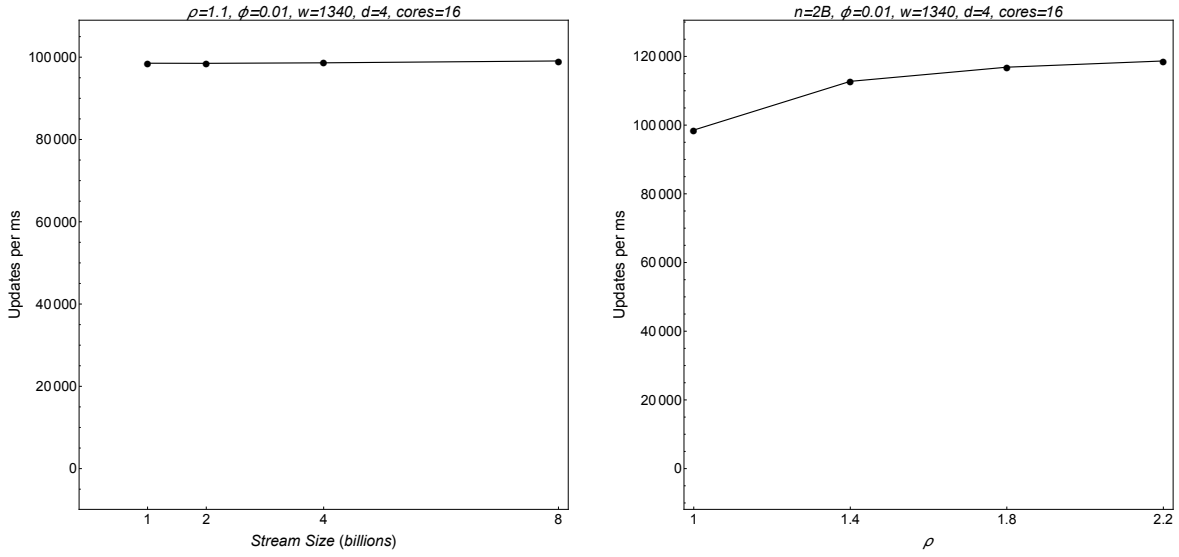

(a) varying n

(b) varying ρ

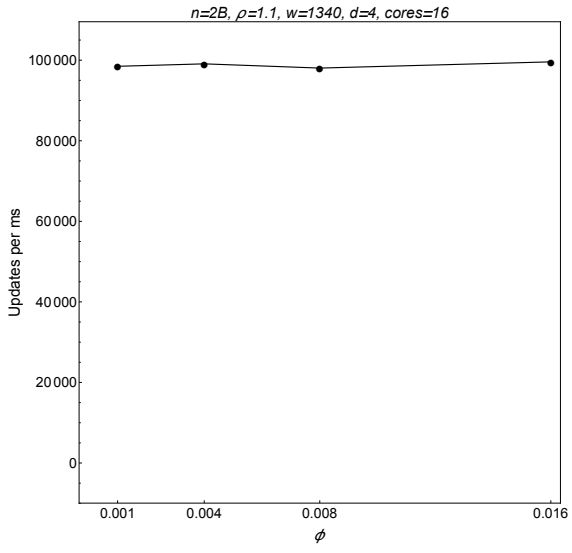

(c) varying ϕ

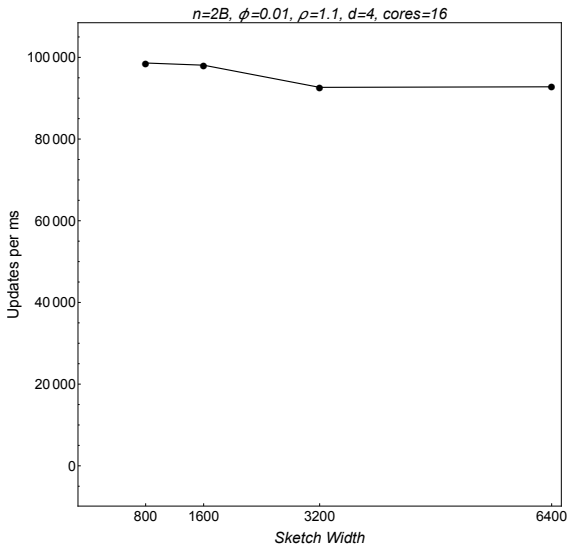

(d) varying w

Figure 4: Updates

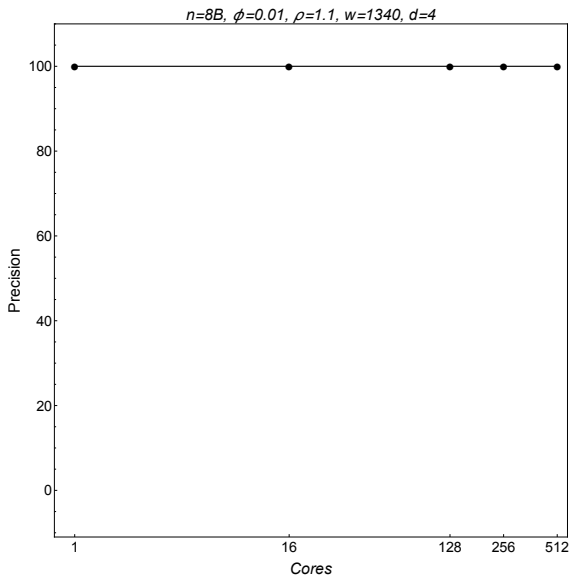

(a) Precision

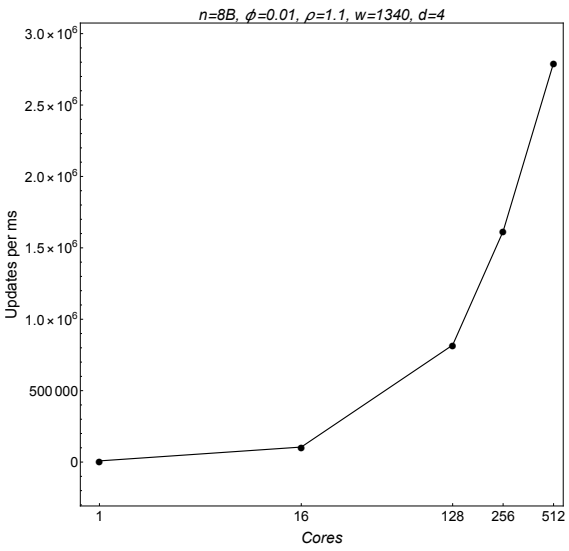

(b) Updates/ms

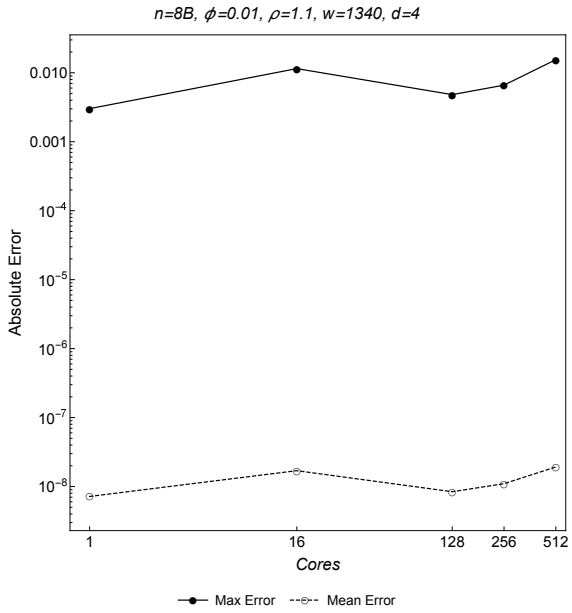

(c) Absolute Error

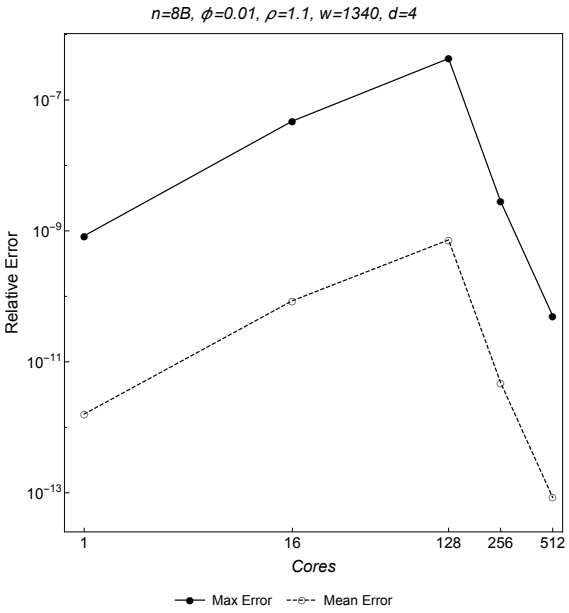

(d) Relative Error

Figure 5: Scalability

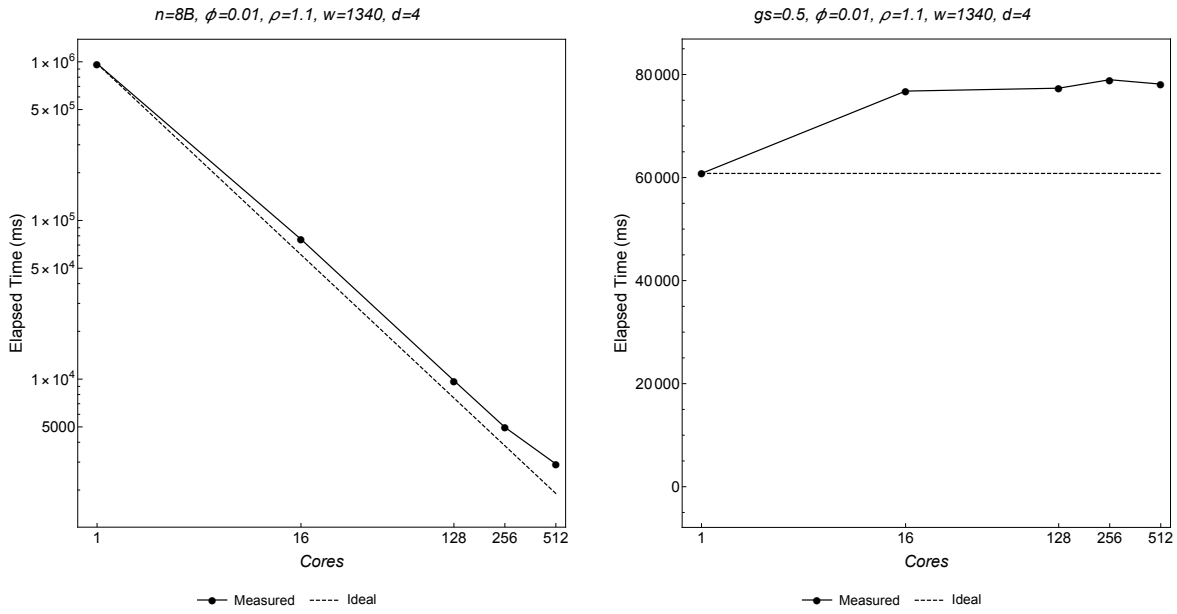

(a) Strong Scalability

(b) Weak Scalability

Figure 6: Parallel Scalability

References

[1] Kevin Beyer and Raghu Ramakrishnan, *Bottom–up computation of sparse and iceberg cubes*, Proceedings of the acm sigmod international conference on management of data. acm, new york, 1999, pp. 359–370.

[2] Sergey Brin, Rajeev Motwani, Jeffrey D. Ullman, and Shalom Tsur, *Dynamic itemset counting and implication rules for market basket data*, Sigmod '97: Proceedings of the 1997 acm sigmod international conference on management of data, 1997, pp. 255–264.

[3] Massimo Cafaro and Marco Pulimeno, *Merging frequent summaries*, Proceedings of the 17th italian conference on theoretical computer science (ictcs 2016), volume 1720, 2016, pp. 280–285.

[4] Massimo Cafaro, Marco Pulimeno, Italo Epicoco, and Giovanni Aloisio, *Mining frequent items in the time fading model*, Information Sciences **370–371** (2016), 221–238.

[5] Massimo Cafaro, Marco Pulimeno, and Piergiulio Tempesta, *A parallel space saving algorithm for frequent items and the hurwitz zeta distribution*, Information Sciences **329** (2016), 1 –19.

[6] Massimo Cafaro and Piergiulio Tempesta, *Finding frequent items in parallel*, Concurr. Comput. : Pract. Exper. **23** (October 2011), no. 15, 1774–1788.

[7] Moses Charikar, Kevin Chen, and Martin Farach-Colton, *Finding frequent items in data streams*, Icalp '02: Proceedings of the 29th international colloquium on automata, languages and programming, 2002, pp. 693–703.

[8] Ling Chen and Qingling Mei, *Mining frequent items in data stream using time fading model*, Information Sciences **257** (2014), 54 –69.

[9] G. Cormode, F. Korn, and S. Tirthapura, *Exponentially decayed aggregates on data streams*, Data engineering, 2008. icde 2008. ieee 24th international conference on, 2008April, pp. 1379–1381.

[10] G. Cormode, V. Shkapenyuk, D. Srivastava, and Bojian Xu, *Forward decay: A practical time decay model for streaming systems*, Data engineering, 2009. icde '09. ieee 25th international conference on, 2009March, pp. 138–149.

[11] Graham Cormode and S. Muthukrishnan, *An improved data stream summary: the count-min sketch and its applications*, J. Algorithms **55** (2005), no. 1, 58–75.

[12] ———, *What's hot and what's not: Tracking most frequent items dynamically*, ACM Trans. Database Syst. **30** (March 2005), no. 1, 249–278.

[13] Sudipto Das, Shyam Antony, Divyakant Agrawal, and Amr El Abbadi, *Thread cooperation in multicore architectures for frequency counting over multiple data streams*, Proc. VLDB Endow. **2** (August 2009), no. 1, 217–228.

[14] Erik D. Demaine, Alejandro López-Ortiz, and J. Ian Munro, *Frequency estimation of internet packet streams with limited space*, Esa, 2002, pp. 348–360.

[15] Ugo Erra and Bernardino Frola, *Frequent items mining acceleration exploiting fast parallel sorting on the {GPU}*, Procedia Computer Science **9** (2012), no. 0, 86 –95. Proceedings of the International Conference on Computational Science, {ICCS} 2012.

[16] Cristian Estan and George Varghese, *New directions in traffic measurement and accounting*, Imw '01: Proceedings of the 1st acm sigcomm workshop on internet measurement, 2001, pp. 75–80.

[17] Min Fang, Narayanan Shivakumar, Hector Garcia-Molina, Rajeev Motwani, and Jeffrey D. Ullman, *Computing iceberg queries efficiently*, Proceedings of the 24th international conference on very large data bases, vldb. morgan–kaufmann, san mateo, calif., 1998, pp. 299–310.

[18] Alexander Gelbukhl (ed.), *Computational linguistics and intelligent text processing, 7th international conference, cicling 2006*, Lecture Notes in Computer Science, vol. 3878, Springer–Verlag, 2006.

[19] Phillip B. Gibbons and Yossi Matias, *Synopsis data structures for massive data sets*, Dimacs: Series in discrete mathematics and theoretical computer science: Special issue on external memory algorithms and visualization, vol. a., 1999, pp. 39–70.

[20] Naga K. Govindaraju, Nikunj Raghuvanshi, and Dinesh Manocha, *Fast and approximate stream mining of quantiles and frequencies using graphics processors*, Proceedings of the 2005 acm sigmod international conference on management of data, 2005, pp. 611–622.

[21] Cheqing Jin, Weining Qian, Chaofeng Sha, Jeffrey X. Yu, and Aoying Zhou, *Dynamically maintaining frequent items over a data stream*, In proc. of cikm, 2003, pp. 287–294.

[22] Richard M. Karp, Scott Shenker, and Christos H. Papadimitriou, *A simple algorithm for finding frequent elements in streams and bags*, ACM Trans. Database Syst. **28** (2003), no. 1, 51–55.

[23] A. Manjhi, V. Shkapenyuk, K. Dhamdhere, and C. Olston, *Finding (recently) frequent items in distributed data streams*, Data engineering, 2005. icde 2005. proceedings. 21st international conference on, 2005April, pp. 767–778.

[24] Gurmeet Singh Manku and Rajeev Motwani, *Approximate frequency counts over data streams*, In vldb, 2002, pp. 346–357.

[25] Ahmed Metwally, Divyakant Agrawal, and Amr El Abbadi, *An integrated efficient solution for computing frequent and top-k elements in data streams*, ACM Trans. Database Syst. **31** (September 2006), no. 3, 1095–1133.

[26] Jayadev Misra and David Gries, *Finding repeated elements*, Sci. Comput. Program. **2** (1982), no. 2, 143–152.

[27] Rong Pan, Lee Breslau, Balaji Prabhakar, and Scott Shenker, *Approximate fairness through differential dropping*, SIGCOMM Comput. Commun. Rev. **33** (2003), no. 2, 23–39.

[28] Pratanu Roy, Jens Teubner, and Gustavo Alonso, *Efficient frequent item counting in multi-core hardware*, Proceedings of the 18th acm sigkdd international conference on knowledge discovery and data mining, 2012, pp. 1451–1459.

[29] Kanat Tangwongsan, Srikanta Tirthapura, and Kun-Lung Wu, *Parallel streaming frequency-based aggregates*, Proceedings of the 26th acm symposium on parallelism in algorithms and architectures, 2014, pp. 236–245.

[30] Yu Zhang, *Parallelizing the weighted lossy counting algorithm in high-speed network monitoring*, Instrumentation, measurement, computer, communication and control (imccc), second international conference on, 2012, pp. 757–761.

[31] Yu Zhang, Yue Sun, Jianzhong Zhang, Jingdong Xu, and Ying Wu, *An efficient framework for parallel and continuous frequent item monitoring*, Concurrency and Computation: Practice and Experience **26** (2014), no. 18, 2856–2879.